arXiv:0908.4144v1 [cs.LG] 28 Aug 2009

ABC-LogitBoost for Multi-class Classification

Ping Li
Department of Statistical Science
Faculty of Computing and Information Science
Cornell University
Ithaca, NY 14853
pingli@cornell.edu

Abstract

We develop *abc-logitboost*, based on the prior work on *abc-boost*[10] and *robust logitboost*[11]. Our extensive experiments on a variety of datasets demonstrate the considerable improvement of *abc-logitboost* over *logitboost* and *abc-mart*.

1 Introduction

Boosting[1] algorithms [14, 4, 5, 2, 15, 7, 13, 6] have become very successful in machine learning. This study revisits *logitboost*[7] under the framework of *adaptive base class boost (abc-boost)* in [10], for multi-class classification.

We denote a training dataset by $\{y_i, \mathbf{x}_i\}_{i=1}^N$, where N is the number of feature vectors (samples), $\mathbf{x}_i$ is the ith feature vector, and $y_i \in \{0, 1, 2, ..., K-1\}$ is the ith class label, where $K \geq 3$ in multi-class classification.

Both *logitboost*[7] and *mart* (multiple additive regression trees)[6] algorithms can be viewed as generalizations to the logistic regression model, which assumes the class probabilities $p_{i,k}$ to be

$$p_{i,k} = \mathbf{Pr}\left(y_i = k|\mathbf{x}_i\right) = \frac{e^{F_{i,k}(\mathbf{x_i})}}{\sum_{s=0}^{K-1} e^{F_{i,s}(\mathbf{x_i})}}, \qquad i = 1, 2, ..., N, \tag{1}$$

While traditional logistic regression assumes $F_{i,k}(\mathbf{x}_i) = \beta^{\mathrm{T}}\mathbf{x}_i$, *logitboost* and *mart* adopt the flexible "additive model," which is a function of M terms:

$$F^{(M)}(\mathbf{x}) = \sum_{m=1}^{M} \rho_m h(\mathbf{x}; \mathbf{a}_m), \tag{2}$$

where $h(\mathbf{x}; \mathbf{a}_m)$, the base learner, is typically a regression tree. The parameters, ρ_m and $\mathbf{a}_m$, are learned from the data, by maximum likelihood, which is equivalent to minimizing the *negative log-likelihood loss*

$$L = \sum_{i=1}^{N} L_i, \qquad L_i = -\sum_{k=0}^{K-1} r_{i,k} \log p_{i,k} \tag{3}$$

where $r_{i,k} = 1$ if $y_i = k$ and $r_{i,k} = 0$ otherwise.

For identifiability, the "sum-to-zero" constraint, $\sum_{k=0}^{K-1} F_{i,k} = 0$, is usually adopted [7, 6, 17, 9, 16, 18].

1.1 Logitboost

As described in Alg. 1, [7] builds the additive model (2) by a greedy stage-wise procedure, using a second-order (diagonal) approximation, which requires knowing the first two derivatives of the loss function (3) with respective

[1]The idea of *abc-logitboost* was included in an unfunded grant proposal submitted in early December 2008.

to the function values $F_{i,k}$. [7] obtained:

$$\frac{\partial L_i}{\partial F_{i,k}} = -\left(r_{i,k} - p_{i,k}\right), \qquad \frac{\partial^2 L_i}{\partial F_{i,k}^2} = p_{i,k}\left(1 - p_{i,k}\right). \tag{4}$$

Those derivatives can be derived by assuming no relations among $F_{i,k}$, $k = 0$ to $K-1$. However, [7] used the "sum-to-zero" constraint $\sum_{k=0}^{K-1} F_{i,k} = 0$ throughout the paper and they provided an alternative explanation. [7] showed (4) by conditioning on a "base class" and noticed the resultant derivatives are independent of the particular choice of the base class.

Algorithm 1 LogitBoost[7, Alg. 6]. ν is the shrinkage (e.g., $\nu = 0.1$).

0: $r_{i,k} = 1$, if $y_i = k$, $r_{i,k} = 0$ otherwise.
1: $F_{i,k} = 0,\ p_{i,k} = \frac{1}{K},\quad k = 0$ to $K-1,\ i = 1$ to N
2: For $m = 1$ to M Do
3: For $k = 0$ to $K-1$, Do
4: Compute $w_{i,k} = p_{i,k}\left(1 - p_{i,k}\right)$.
5: Compute $z_{i,k} = \frac{r_{i,k} - p_{i,k}}{p_{i,k}\left(1 - p_{i,k}\right)}$.
6: Fit the function $f_{i,k}$ by a weighted least-square of $z_{i,k}$ to $\mathbf{x}_i$ with weights $w_{i,k}$.
7: $F_{i,k} = F_{i,k} + \nu\frac{K-1}{K}\left(f_{i,k} - \frac{1}{K}\sum_{k=0}^{K-1} f_{i,k}\right)$
8: End
9: $p_{i,k} = \exp(F_{i,k}) / \sum_{s=0}^{K-1}\exp(F_{i,s}),\quad k = 0$ to $K-1,\ i = 1$ to N
10: End

At each stage, *logitboost* fits an individual regression function separately for each class. This is analogous to the popular *individualized regression* approach in multinomial logistic regression, which is known [3, 1] to result in loss of statistical efficiency, compared to the full (conditional) maximum likelihood approach.

On the other hand, in order to use trees as base learner, the diagonal approximation appears to be a must, at least from the practical perspective.

1.2 Adaptive Base Class Boost

[10] derived the derivatives of (3) under the sum-to-zero constraint. Without loss of generality, we can assume that class 0 is the base class. For any $k \neq 0$,

$$\frac{\partial L_i}{\partial F_{i,k}} = \left(r_{i,0} - p_{i,0}\right) - \left(r_{i,k} - p_{i,k}\right), \qquad \frac{\partial^2 L_i}{\partial F_{i,k}^2} = p_{i,0}(1 - p_{i,0}) + p_{i,k}(1 - p_{i,k}) + 2p_{i,0}p_{i,k}. \tag{5}$$

The base class must be identified at each boosting iteration during training. [10] suggested an exhaustive procedure to adaptively find the best base class to minimize the training loss (3) at each iteration.

[10] combined the idea of *abc-boost* with *mart*. The algorithm, *abc-mart*, achieved good performance in multi-class classification on the datasets used in [10].

1.3 Our Contributions

We propose *abc-logitboost*, by combining *abc-boost* with *robust logitboost*[11]. Our extensive experiments will demonstrate that *abc-logitboost* can considerably improve *logitboost* and *abc-mart* on a variety of datasets.

2 Robust Logitboost

Our work is based on *robust logitboost*[11], which differs from the original *logitboost* algorithm. Thus, this section provides an introduction to *robust logitboost*.

[6, 8] commented that *logitboost* (Alg. 1) can be numerically unstable. The original paper[7] suggested some "crucial implementation protections" on page 17 of [7]:

- In Line 5 of Alg. 1, compute the response $z_{i,k}$ by $\frac{1}{p_{i,k}}$ (if $r_{i,k}=1$) or $\frac{-1}{1-p_{i,k}}$ (if $r_{i,k}=0$).
- Bound the response $|z_{i,k}|$ by $z_{max} \in [2,4]$.

Note that the above operations are applied to each individual sample. The goal is to ensure that the response $|z_{i,k}|$ is not too large (Note that $|z_{i,k}| > 1$ always). On the other hand, we should hope to use larger $|z_{i,k}|$ to better capture the data variation. Therefore, the thresholding occurs very frequently and it is expected that some of the useful information is lost.

[11] demonstrated that, if implemented carefully, *logitboost* is almost identical to *mart*. The only difference is the tree-splitting criterion.

2.1 The Tree-Splitting Criterion Using the Second-Order Information

Consider N weights w_i, and N response values z_i, $i=1$ to N, which are assumed to be ordered according to the sorted order of the corresponding feature values. The tree-splitting procedure is to find the index s, $1 \le s < N$, such that the weighted mean square error (MSE) is reduced the most if split at s. That is, we seek s to maximize

$$Gain(s) = MSE_T - (MSE_L + MSE_R) = \sum_{i=1}^{N}(z_i-\bar{z})^2 w_i - \left[\sum_{i=1}^{s}(z_i-\bar{z}_L)^2 w_i + \sum_{i=s+1}^{N}(z_i-\bar{z}_R)^2 w_i\right]$$

where $\bar{z} = \frac{\sum_{i=1}^{N} z_i w_i}{\sum_{i=1}^{N} w_i}$, $\bar{z}_L = \frac{\sum_{i=1}^{s} z_i w_i}{\sum_{i=1}^{s} w_i}$, and $\bar{z}_R = \frac{\sum_{i=s+1}^{N} z_i w_i}{\sum_{i=s+1}^{N} w_i}$. After simplification, we obtain

$$Gain(s) = \frac{\left[\sum_{i=1}^{s} z_i w_i\right]^2}{\sum_{i=1}^{s} w_i} + \frac{\left[\sum_{i=s+1}^{N} z_i w_i\right]^2}{\sum_{i=s+1}^{N} w_i} - \frac{\left[\sum_{i=1}^{N} z_i w_i\right]^2}{\sum_{i=1}^{N} w_i}$$

Plugging in $w_i = p_{i,k}(1-p_{i,k})$, and $z_i = \frac{r_{i,k}-p_{i,k}}{p_{i,k}(1-p_{i,k})}$ as in Alg. 1, yields,

$$Gain(s) = \frac{\left[\sum_{i=1}^{s} r_{i,k}-p_{i,k}\right]^2}{\sum_{i=1}^{s} p_{i,k}(1-p_{i,k})} + \frac{\left[\sum_{i=s+1}^{N} r_{i,k}-p_{i,k}\right]^2}{\sum_{i=s+1}^{N} p_{i,k}(1-p_{i,k})} - \frac{\left[\sum_{i=1}^{N} r_{i,k}-p_{i,k}\right]^2}{\sum_{i=1}^{N} p_{i,k}(1-p_{i,k})}.$$

Because the computations involve $\sum p_{i,k}(1-p_{i,k})$ as a group, this procedure is actually numerically stable.

In comparison, *mart*[6] only used the first order information to construct the trees, i.e.,

$$MARTGain(s) = \left[\sum_{i=1}^{s} r_{i,k}-p_{i,k}\right]^2 + \left[\sum_{i=s+1}^{N} r_{i,k}-p_{i,k}\right]^2 - \left[\sum_{i=1}^{N} r_{i,k}-p_{i,k}\right]^2.$$

2.2 The Robust Logitboost Algorithm

Algorithm 2 Robust logitboost, which is very similar to mart, except for Line 4.

1: $F_{i,k}=0$, $p_{i,k}=\frac{1}{K}$, $k=0$ to $K-1$, $i=1$ to N
2: For $m=1$ to M Do
3: For $k=0$ to $K-1$ Do
4: $\{R_{j,k,m}\}_{j=1}^{J} = J$-terminal node regression tree from $\{r_{i,k}-p_{i,k},\ \mathbf{x}_i\}_{i=1}^{N}$,
: with weights $p_{i,k}(1-p_{i,k})$ as in Section 2.1.
5: $\beta_{j,k,m} = \frac{K-1}{K}\frac{\sum_{\mathbf{x}_i \in R_{j,k,m}} r_{i,k}-p_{i,k}}{\sum_{\mathbf{x}_i \in R_{j,k,m}} (1-p_{i,k})p_{i,k}}$
6: $F_{i,k} = F_{i,k} + \nu \sum_{j=1}^{J} \beta_{j,k,m} 1_{\mathbf{x}_i \in R_{j,k,m}}$
7: End
8: $p_{i,k} = \exp(F_{i,k}) / \sum_{s=0}^{K-1} \exp(F_{i,s}), \quad k=0$ to $K-1,\ i=1$ to N
9: End

Alg. 2 describes *robust logitboost* using the tree-splitting criterion developed in Section 2.1. Note that after trees are constructed, the values of the terminal nodes are computed by

$$\frac{\sum_{node} z_{i,k} w_{i,k}}{\sum_{node} w_{i,k}} = \frac{\sum_{node} r_{i,k} - p_{i,k}}{\sum_{node} p_{i,k}(1-p_{i,k})},$$

which explains Line 5 of Alg. 2.

2.2.1 Three Main Parameters: J, ν, and M

Alg. 2 has three main parameters, to which the performance is not very sensitive, as long as they fall in some reasonable range. This is a very significant advantage in practice.

The number of terminal nodes, J, determines the capacity of the base learner. [6] suggested $J = 6$. [7, 18] commented that $J > 10$ is unlikely. In our experience, for large datasets (or moderate datasets in high-dimensions), $J = 20$ is often a reasonable choice; also see [12].

The shrinkage, ν, should be large enough to make sufficient progress at each step and small enough to avoid over-fitting. [6] suggested $\nu \leq 0.1$. Normally, $\nu = 0.1$ is used.

The number of iterations, M, is largely determined by the affordable computing time. A commonly-regarded merit of boosting is that over-fitting can be largely avoided for reasonable J and ν.

3 Adaptive Base Class Logitboost

Algorithm 3 *Abc-logitboost* using the exhaustive search strategy for the base class, as suggested in [10]. The vector B stores the base class numbers.

1: $F_{i,k} = 0,\ p_{i,k} = \frac{1}{K},\quad k = 0$ to $K-1,\ i = 1$ to N
2: For $m = 1$ to M Do
3: For $b = 0$ to $K-1$, Do
4: For $k = 0$ to $K-1$, $k \neq b$, Do
5: $\{R_{j,k,m}\}_{j=1}^{J} = J$-terminal node regression tree from $\{-(r_{i,b} - p_{i,b}) + (r_{i,k} - p_{i,k}),\ \mathbf{x}_i\}_{i=1}^{N}$
: with weights $p_{i,b}(1-p_{i,b}) + p_{i,k}(1-p_{i,k}) + 2p_{i,b}p_{i,k}$, as in Section 2.1.
6: $\beta_{j,k,m} = \frac{\sum_{\mathbf{x}_i \in R_{j,k,m}} -(r_{i,b} - p_{i,b}) + (r_{i,k} - p_{i,k})}{\sum_{\mathbf{x}_i \in R_{j,k,m}} p_{i,b}(1-p_{i,b}) + p_{i,k}(1-p_{i,k}) + 2p_{i,b}p_{i,k}}$
7: $G_{i,k,b} = F_{i,k} + \nu \sum_{j=1}^{J} \beta_{j,k,m} 1_{\mathbf{x}_i \in R_{j,k,m}}$
8: End
9: $G_{i,b,b} = -\sum_{k \neq b} G_{i,k,b}$
10: $q_{i,k} = \exp(G_{i,k,b}) / \sum_{s=0}^{K-1} \exp(G_{i,s,b})$
11: $L^{(b)} = -\sum_{i=1}^{N} \sum_{k=0}^{K-1} r_{i,k} \log(q_{i,k})$
12: End
13: $B(m) = \underset{b}{\text{argmin}}\ L^{(b)}$
14: $F_{i,k} = G_{i,k,B(m)}$
15: $p_{i,k} = \exp(F_{i,k}) / \sum_{s=0}^{K-1} \exp(F_{i,s})$
16: End

The recently proposed *abc-boost* [10] algorithm consists of two key components:

1. Using the widely-used *sum-to-zero* constraint[7, 6, 17, 9, 16, 18] on the loss function, one can formulate boosting algorithms only for $K-1$ classes, by using one class as the **base class**.

2. At each boosting iteration, **adaptively** select the base class according to the training loss. [10] suggested an exhaustive search strategy.

[10] combined *abc-boost* with *mart* to develop *abc-mart*. [10] demonstrated the good performance of *abc-mart* compared to *mart*. This study will illustrate that ***abc-logitboost***, the combination of *abc-boost* with *(robust) logitboost*, will further reduce the test errors, at least on a variety of datasets.

Alg. 3 presents *abc-logitboost*, using the derivatives in (5) and the same exhaustive search strategy as in *abc-mart*. Again, *abc-logitboost* differs from *abc-mart* only in the tree-splitting procedure (Line 5 in Alg. 3).

4 Experiments

Table 1 lists the datasets in our experiments, which include all the datasets used in [10], plus *Mnist10k*[2].

Table 1: For *Letter*, *Pendigits*, *Zipcode*, *Optdigits* and *Isolet*, we used the standard (default) training and test sets. For *Covertype*, we use the same split in [10]. For *Mnist10k*, we used the original 10000 test samples in the original *Mnist* dataset for training, and the original 60000 training samples for testing. Also, as explained in [10], *Letter2k* (*Letter4k*) used the last 2000 (4000) samples of *Letter* for training and the remaining 18000 (16000) for testing, from the original *Letter* dataset.

dataset	K	# training	# test	# features
Covertype	7	290506	290506	54
Mnist10k	10	10000	60000	784
Letter2k	26	2000	18000	16
Letter4k	26	4000	16000	16
Letter	26	16000	4000	16
Pendigits	10	7494	3498	16
Zipcode	10	7291	2007	256
Optdigits	10	3823	1797	64
Isolet	26	6218	1559	617

Note that *Zipcode*, *Otpdigits*, and *Isolet* are very small datasets (especially the testing sets). They may not necessarily provide a reliable comparison of different algorithms. Since they are popular datasets, we nevertheless include them in our experiments.

Recall *logitboost* has three main parameters, J, ν, and M. Since overfitting is largely avoided, we simply let $M = 10000$ ($M = 5000$ only for *Covertype*), unless the machine zero is reached. The performance is not sensitive to ν (as long as $\nu \leq 0.1$). The performance is also not too sensitive to J in a good range.

Ideally, we would like to show that, for every reasonable combination of J and ν (using M as large as possible), *abc-logitboost* exhibits consistent improvement over *(robust) logitboost*. For most datasets, we experimented with every combination of $J \in \{4, 6, 8, 10, 12, 14, 16, 18, 20\}$ and $\nu \in \{0.04, 0.06, 0.08, 0.1\}$.

We provide a summary of the experiments after presenting the detailed results on *Mnist10k*.

4.1 Experiments on the *Mnist10k* Dataset

For this dataset, we experimented with every combination of $J \in \{4, 6, 8, 10, 12, 14, 16, 18, 20\}$ and $\nu \in \{0.04, 0.06, 0.08, 0.1\}$. We trained till the loss (3) reached the machine zero, to exhaust the capacity of the learner so that we could provide a reliable comparison, up to $M = 10000$ iterations.

Figures 1 and 2 present the mis-classification errors for every ν, J, and M:

- Essentially no ovefitting is observed, especially for *abc-logitboost*. This is why we simply report the smallest test error in Table 2.
- The performance is not sensitive to ν.
- The performance is not very sensitive to J, for $J = 8$ to 20.

Interestingly, *abc-logitboost* sometimes needed more iterations to reach machine zero than *logitboost*. This can be explained in part by the fact that the "ν" in *logitboost* is not precisely the same "ν" in *abc-logitboost*[10]. This is also why we would like to experiment with a range of ν values.

[2]We also did limited experiments on the original *Mnist* dataset (i.e., 60000 training samples and 10000 testing samples), the test mis-classification error rate was about 1.3%.

Table 2 summarizes the smallest test mis-classification errors along with the relative improvements (denoted by R_{err}) of *abc-logitboost* over *logitboost*. For most J and ν, *abc-logitboost* exhibits about $R_{err} = 12 \sim 15(\%)$ smaller test mis-classification errors than *logitboost*. The P-values range from 1.9×10^{-10} to 3.9×10^{-5}, although they are not reported in Table 2.

Table 2: ***Mnist10k***. The test mis-classification errors of *logitboost* and ***abc-logitboost***, along with the relative improvement R_{err} (%). For each J and ν, we report the smallest values in Figures 1 and 2. Each cell contains three numbers, which are *logitboost error*, ***abc-logitboost error***, and relative improvement R_{err} (%).

	$\nu = 0.04$	$\nu = 0.06$	$\nu = 0.08$	$\nu = 0.1$
$J = 4$	2911 **2623** 9.9	2884 **2597** 10.0	2876 **2530** 12.0	2878 **2485** 13.7
$J = 6$	2658 **2255** 15.2	2644 **2240** 15.3	2625 **2224** 15.3	2626 **2212** 15.8
$J = 8$	2536 **2157** 14.9	2541 **2122** 16.5	2521 **2117** 16.0	2533 **2134** 15.8
$J = 10$	2486 **2118** 14.8	2472 **2111** 14.6	2447 **2083** 14.9	2446 **2095** 14.4
$J = 12$	2435 **2082** 14.5	2424 **2086** 13.9	2420 **2086** 13.8	2426 **2090** 13.9
$J = 14$	2399 **2083** 13.2	2407 **2081** 13.5	2402 **2056** 14.4	2400 **2048** 14.7
$J = 16$	2421 **2098** 13.3	2405 **2114** 12.1	2382 **2083** 12.6	2364 **2079** 12.1
$J = 18$	2397 **2086** 13.0	2397 **2079** 13.3	2386 **2080** 12.8	2357 **2085** 11.5
$J = 20$	2384 **2124** 10.9	2409 **2109** 14.5	2404 **2095** 12.9	2372 **2101** 11.4

The original *abc-boost* paper[10] did not include experiments on *Mnist10k*. Thus, in this study, Table 3 summarizes the smallest test mis-classification errors for *mart* and *abc-mart*. Again, we can see very consistent and considerable improvement of *abc-mart* over *mart*. Also, comparing Tables 2 and 3, we can see that *abc-logitboost* also significantly improves *abc-mart*.

Table 3: ***Mnist10k***. The test mis-classification errors of *mart* and ***abc-mart***, along with the relative improvement R_{err} (%). For each J and ν, we report the smallest values in Figures 1 and 2. Each cell contains three numbers, which are *mart error*, ***abc-mart error***, and relative improvement R_{err} (%).

	$\nu = 0.04$	$\nu = 0.06$	$\nu = 0.08$	$\nu = 0.1$
$J = 4$	3346 **3054** 8.7	3308 **3009** 9.0	3302 **2855** 13.5	3287 **2792** 15.1
$J = 6$	3176 **2752** 13.4	3074 **2624** 14.6	3071 **2649** 13.7	3089 **2572** 16.7
$J = 8$	3040 **2557** 15.9	3012 **2552** 15.2	3000 **2529** 15.7	2993 **2566** 14.3
$J = 10$	2979 **2537** 14.8	2941 **2515** 14.5	2957 **2509** 15.2	2947 **2493** 15.4
$J = 12$	2912 **2498** 14.2	2897 **2453** 15.3	2906 **2475** 14.8	2887 **2469** 14.5
$J = 14$	2907 **2473** 14.9	2886 **2466** 14.6	2874 **2463** 14.3	2864 **2435** 15.0
$J = 16$	2885 **2466** 14.5	2879 **2441** 15.2	2868 **2459** 14.2	2854 **2451** 14.1
$J = 18$	2852 **2467** 13.5	2860 **2447** 14.4	2865 **2436** 15.0	2852 **2448** 14.2
$J = 20$	2831 **2438** 13.9	2833 **2440** 13.9	2832 **2425** 14.4	2813 **2434** 13.5

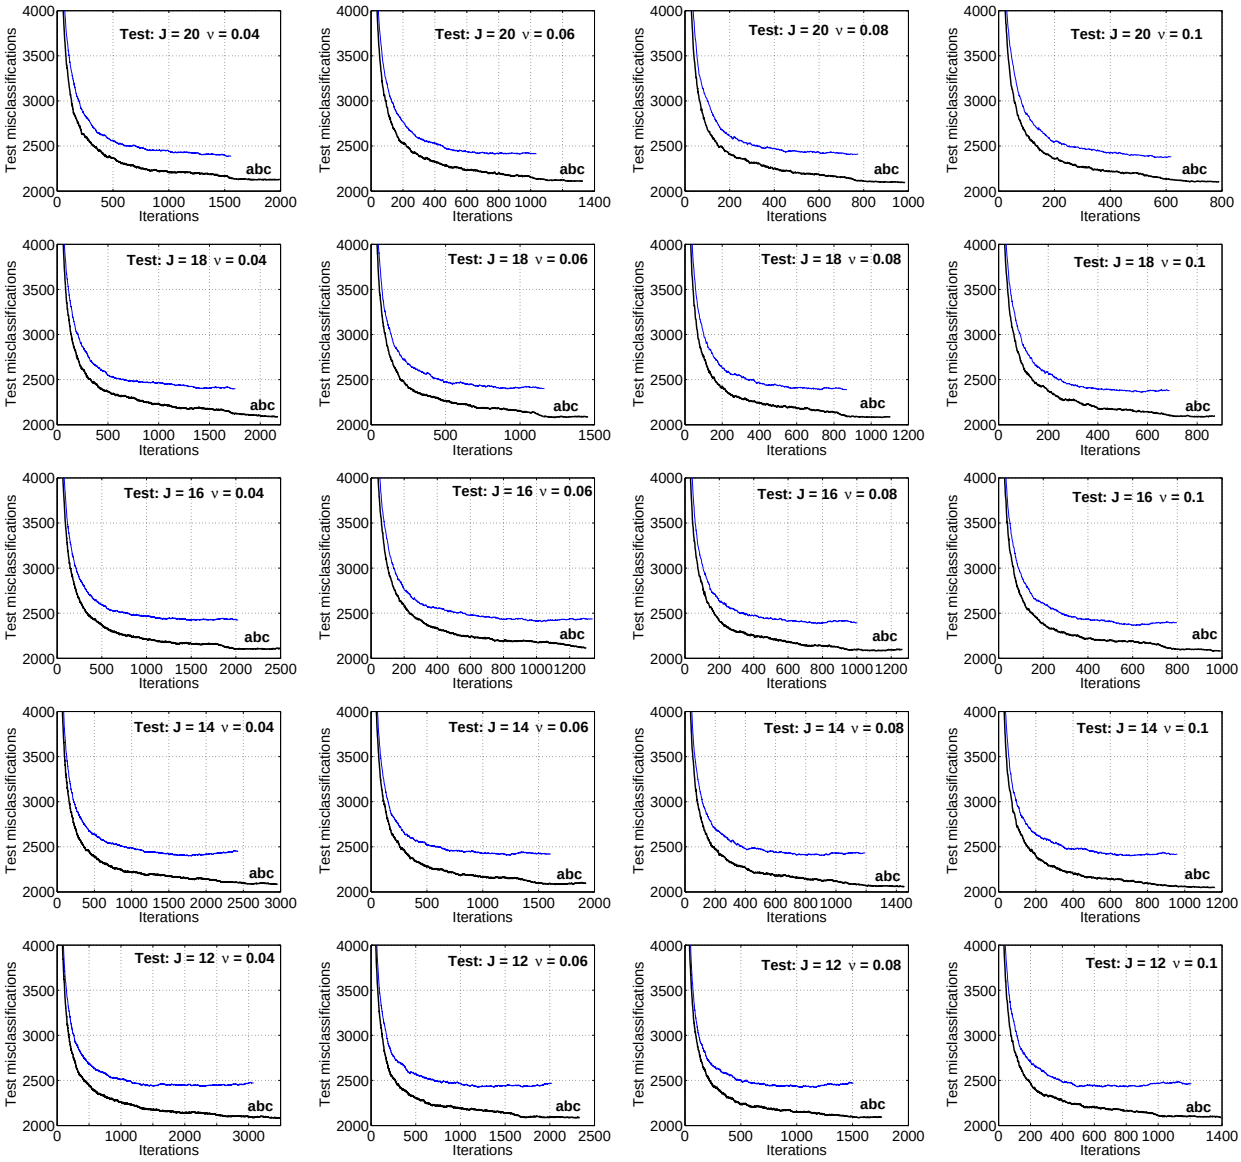

Figure 1: ***Mnist10k***. The test mis-classification errors, for *logitboost* and *abc-logitboost*. $J = 12$ to 20.

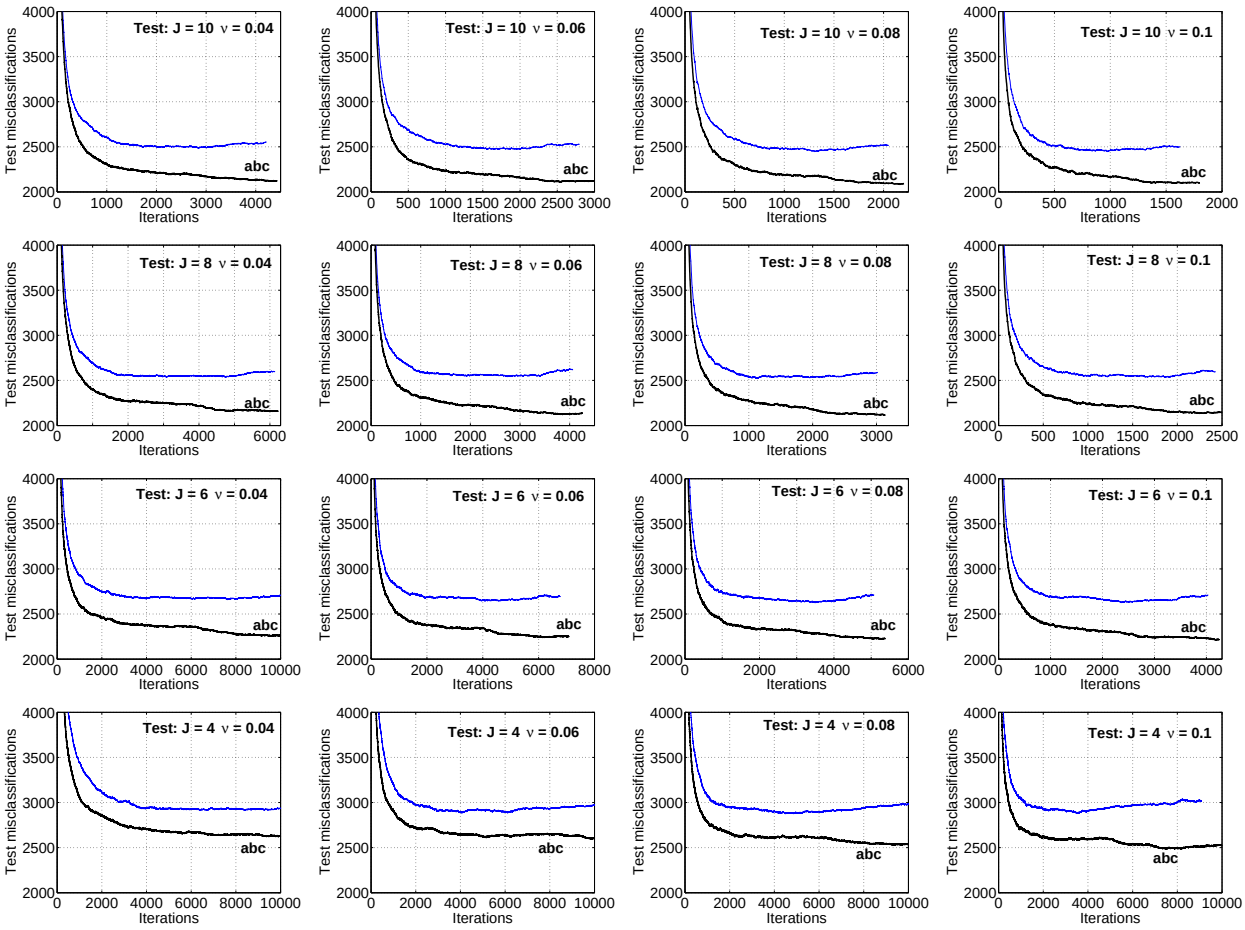

Figure 2: ***Mnist10k***. The test mis-classification errors, for *logitboost* and *abc-logitboost*. $J = 4$ to 10.

4.2 Summary of Test Mis-Classification Errors

Table 4 summarizes the test errors, which are the overall best (smallest) test mis-classification errors. In the table, R_{err} (%) is the relative improvement of test performance. The P-values tested the statistical significance if *abc-logitboost* achieved smaller **error rates** than *logitboost*.

To compare *abc-logitboost* with *abc-mart*, Table 4 also includes the test errors for *abc-mart* and the P-values (i.e., P-value (2)) for testing the statistical significance if *abc-logitboost* achieved smaller **error rates** than *abc-mart*. The comparisons indicate that there is a clear performance gap between *abc-logitboost* and *abc-mart*, especially on the large datasets.

Table 4: Summary of test mis-classification errors.

Dataset	**logit**	**abc-logit**	R_{err} (%)	P-value	**abc-mart**	P-vlaue (2)
Covertype	10759	9693	9.9	1.6×10^{-14}	10375	4.8×10^{-7}
Mnist10k	2357	2048	13.1	1.0×10^{-6}	2425	4.6×10^{-9}
Letter2k	2257	1984	12.1	4.0×10^{-6}	2180	6.2×10^{-4}
Letter4k	1220	1031	15.5	1.8×10^{-5}	1126	0.017
Letter	107	89	16.8	9.7×10^{-3}	99	0.23
Pendigits	109	90	17.4	8.6×10^{-3}	100	0.23
Zipcode	103	92	10.7	0.21	100	0.28
Optdigits	49	38	22.5	0.11	43	0.29
Isolet	62	55	11.3	0.25	64	0.20

4.3 Experiments on the *Covertype* Dataset

Table 5 summarizes the smallest test mis-classification errors of *logitboost* and *abc-logitboost*, along with the relative improvements (R_{err}). Since this is a fairly large dataset, we only experimented with $\nu = 0.1$ and $J = 10$ and 20.

Table 5: ***Covertype***. We report the test mis-classification errors of *logitboost* and *abc-logitboost*, together with the relative improvements (R_{err}, $\%$) in parentheses.

ν	M	J	logit	abc-logit
0.1	1000	10	29865	23774 (20.4)
0.1	1000	20	19443	14443 (25.7)
0.1	2000	10	21620	16991 (21.4)
0.1	2000	20	13914	11336 (18.5)
0.1	3000	10	17805	14295 (19.7)
0.1	3000	20	12076	10399 (13.9)
0.1	5000	10	14698	12185 (17.1)
0.1	5000	20	10759	9693 (9.9)

The results on *Covertype* are reported differently from other datasets. *Covertype* is fairly large. Building a very large model (e.g., $M = 5000$ boosting steps) would be expensive. Testing a very large model at run-time can be costly or infeasible for certain applications (e.g., search engines). Therefore, it is often important to examine the performance of the algorithm at much earlier boosting iterations. Table 5 shows that *abc-logitboost* may improve *logitboost* as much as $R_{err} \approx 20\%$, as opposed to the reported $R_{err} = 9.9\%$ in Table 4.

4.4 Experiments on the *Letter2k* Dataset

Table 6: ***Letter2k***. The test mis-classification errors of *logitboost* and ***abc-logitboost***, along with the relative improvement R_{err} (%). Each cell contains three numbers, which are *logitboost error*, ***abc-logitboost error***, and relative improvement R_{err} (%).

	$\nu = 0.04$	$\nu = 0.06$	$\nu = 0.08$	$\nu = 0.1$
$J = 4$	2576 **2317** 10.1	2535 **2294** 9.5	2545 **2252** 11.5	2523 **2224** 11.9
$J = 6$	2389 **2133** 10.7	2391 **2111** 11.7	2376 **2070** 12.9	2370 **2064** 12.9
$J = 8$	2325 **2074** 10.8	2299 **2046** 11.0	2298 **2033** 11.5	2271 **2025** 10.8
$J = 10$	2294 **2041** 11.0	2292 **1995** 13.0	2279 **2018** 11.5	2276 **2000** 12.1
$J = 12$	2314 **2010** 13.1	2304 **1990** 13.6	2311 **2010** 13.0	2268 **2018** 11.0
$J = 14$	2315 **2015** 13.0	2300 **2003** 12.9	2312 **2003** 13.4	2277 **2024** 11.1
$J = 16$	2302 **2022** 12.2	2394 **1996** 13.0	2276 **3002** 12.0	2257 **1997** 11.5
$J = 18$	2295 **2041** 11.1	2275 **2021** 11.2	2301 **1984** 13.8	2281 **2020** 11.4
$J = 20$	2280 **2047** 10.2	2267 **2020** 10.9	2294 **2020** 11.9	2306 **2031** 11.9

4.5 Experiments on the *Letter4k* Dataset

Table 7: ***Letter4k***. The test mis-classification errors of *logitboost* and ***abc-logitboost***, along with the relative improvement R_{err} (%).

	$\nu = 0.04$	$\nu = 0.06$	$\nu = 0.08$	$\nu = 0.1$
$J = 4$	1460 **1295** 11.3	1471 **1232** 16.2	1452 **1199** 17.4	1446 **1204** 16.7
$J = 6$	1390 **1135** 18.3	1394 **1116** 20.0	1382 **1088** 21.3	1374 **1070** 22.1
$J = 8$	1336 **1078** 19.3	1332 **1074** 19.4	1311 **1062** 19.0	1297 **1042** 20.0
$J = 10$	1289 **1051** 18.5	1285 **1065** 17.1	1280 **1031** 19.5	1273 **1046** 17.8
$J = 12$	1251 **1055** 15.7	1247 **1065** 14.6	1261 **1044** 17.2	1243 **1051** 15.4
$J = 14$	1247 **1060** 15.0	1233 **1050** 14.8	1251 **1037** 17.1	1244 **1060** 14.8
$J = 16$	1244 **1070** 14.0	1227 **1064** 13.3	1231 **1044** 15.2	1228 **1038** 15.5
$J = 18$	1243 **1057** 15.0	1250 **1037** 17.0	1234 **1049** 15.0	1220 **1055** 13.5
$J = 20$	1226 **1078** 12.0	1242 **1069** 13.9	1242 **1054** 15.1	1235 **1051** 14.9

4.6 Experiments on the *Letter* Dataset

Table 8: ***Letter***. The test mis-classification errors of *logitboost* and ***abc-logitboost***, along with the relative improvement R_{err} (%).

	$\nu = 0.04$	$\nu = 0.06$	$\nu = 0.08$	$\nu = 0.1$
$J = 4$	149 **125** 16.1	151 **121** 19.9	148 **122** 17.6	149 **119** 20.1
$J = 6$	130 **112** 13.8	132 **107** 18.9	133 **101** 24.1	129 **102** 20.9
$J = 8$	129 **104** 19.4	125 **102** 18.4	131 **93** 29.0	113 **95** 15.9
$J = 10$	114 **101** 11.4	115 **100** 13.0	123 **96** 22.0	117 **93** 20.5
$J = 12$	112 **96** 14.3	115 **100** 13.0	107 **95** 11.2	112 **95** 15.2
$J = 14$	110 **96** 12.7	113 **98** 13.3	113 **94** 16.8	110 **89** 19.1
$J = 16$	111 **97** 12.6	113 **94** 16.8	109 **93** 14.7	109 **95** 12.8
$J = 18$	114 **95** 16.7	112 **92** 17.9	111 **96** 13.5	117 **93** 20.5
$J = 20$	113 **95** 15.9	113 **97** 14.2	115 **93** 19.1	113 **89** 21.2

4.7 Experiments on the *Pendigits* Dataset

Table 9: ***Pendigits***. The test mis-classification errors of *logitboost* and ***abc-logitboost***, along with the relative improvement R_{err} (%).

	$\nu = 0.04$	$\nu = 0.06$	$\nu = 0.08$	$\nu = 0.1$
$J = 4$	119 **92** 22.7	120 **93** 22.5	118 **90** 23.7	119 **92** 22.7
$J = 6$	111 **98** 11.7	111 **97** 12.6	111 **96** 13.5	107 **93** 13.1
$J = 8$	116 **97** 16.4	117 **94** 19.7	115 **95** 17.4	114 **93** 18.4
$J = 10$	116 **100** 13.8	115 **98** 14.8	116 **97** 16.4	111 **97** 12.6
$J = 12$	117 **98** 16.2	113 **98** 13.2	113 **98** 13.3	114 **98** 14.0
$J = 14$	113 **100** 11.5	115 **101** 12.2	112 **99** 11.6	114 **98** 14.0
$J = 16$	112 **100** 10.7	118 **97** 18.8	112 **98** 12.5	113 **96** 15.0
$J = 18$	114 **102** 10.5	112 **97** 13.4	109 **99** 9.2	112 **97** 13.4
$J = 20$	112 **106** 5.4	116 **102** 12.1	113 **100** 11.5	107 **100** 6.5

4.8 Experiments on the *Zipcode* Dataset

Table 10: ***Zipcode***. The test mis-classification errors of *logitboost* and ***abc-logitboost***, along with the relative improvement R_{err} (%).

	$\nu = 0.04$	$\nu = 0.06$	$\nu = 0.08$	$\nu = 0.1$
$J = 4$	114 **111** 2.6	117 **108** 7.6	111 **114** -2.7	115 **107** 7.0
$J = 6$	109 **101** 7.3	107 **102** 4.6	106 **98** 7.5	110 **99** 10.0
$J = 8$	110 **99** 10.0	108 **95** 12.0	108 **96** 11.1	108 **98** 9.3
$J = 10$	111 **97** 12.6	110 **94** 14.5	106 **97** 8.5	103 **94** 8.7
$J = 12$	111 **98** 11.7	112 **98** 12.5	111 **99** 10.8	108 **93** 13.9
$J = 14$	112 **100** 10.7	108 **99** 8.3	110 **97** 11.8	114 **92** 19.3
$J = 16$	111 **98** 11.7	114 **95** 16.7	110 **99** 10.0	111 **98** 11.7
$J = 18$	112 **96** 14.2	114 **98** 14.0	109 **101** 7.3	113 **98** 13.3
$J = 20$	114 **97** 14.9	108 **96** 11.1	109 **100** 8.3	116 **96** 17.2

4.9 Experiments on the *Optdigits* Dataset

Table 11: ***Optdigits***.The test mis-classification errors of *logitboost* and ***abc-logitboost***, along with the relative improvement R_{err} (%).

	$\nu = 0.04$	$\nu = 0.06$	$\nu = 0.08$	$\nu = 0.1$
$J = 4$	52 **41** 21.2	50 **42** 16.0	50 **40** 20.0	49 **41** 16.3
$J = 6$	52 **43** 17.3	52 **45** 13.5	53 **44** 17.0	52 **38** 26.9
$J = 8$	55 **44** 20.0	55 **44** 20.0	53 **45** 15.1	54 **45** 16.7
$J = 10$	57 **50** 12.3	56 **50** 10.7	54 **46** 14.8	55 **42** 23.6
$J = 12$	52 **50** 3.8	55 **48** 12.7	54 **47** 13.0	54 **46** 14.8
$J = 14$	58 **48** 17.2	55 **46** 16.4	56 **51** 8.9	53 **48** 9.4
$J = 16$	61 **54** 11.5	57 **51** 10.5	58 **49** 15.5	56 **46** 17.9
$J = 18$	65 **54** 16.9	64 **55** 14.0	60 **53** 11.7	66 **51** 22.7
$J = 20$	63 **61** 3.2	61 **56** 8.2	64 **55** 14.1	64 **55** 14.1

4.10 Experiments on the *Isolet* Dataset

For this dataset, [10] only experimented with $\nu = 0.1$ for *mart* and *abc-mart*. We add the experiment results for $\nu = 0.06$.

Table 12: ***Isolet***. The test mis-classification errors of *logitboost* and ***abc-logitboost***, along with the relative improvement R_{err} (%).

	$\nu = 0.06$	$\nu = 0.1$
$J = 4$	65 **55** 15.4	62 **55** 11.3
$J = 6$	67 **59** 11.9	69 **58** 15.9
$J = 8$	72 **57** 20.8	72 **60** 16.7
$J = 10$	73 **61** 16.4	75 **62** 17.3
$J = 12$	75 **63** 16.0	75 **64** 14.7
$J = 14$	74 **65** 12.2	75 **60** 20.0
$J = 16$	70 **64** 8.6	71 **62** 12.7
$J = 18$	74 **67** 9.5	73 **62** 15.1
$J = 20$	71 **63** 11.3	73 **65** 11.0

Table 13: ***Isolet***. The test mis-classification errors of *mart* and ***abc-mart***, along with the relative improvement R_{err} (%).

	$\nu = 0.06$	$\nu = 0.1$
$J = 4$	81 **68** 16.1	80 **64** 20.0
$J = 6$	86 **71** 17.4	84 **67** 20.2
$J = 8$	86 **72** 16.3	84 **72** 14.3
$J = 10$	87 **74** 14.9	82 **74** 9.8
$J = 12$	93 **73** 21.5	91 **74** 18.7
$J = 14$	92 **73** 20.7	95 **74** 22.1
$J = 16$	91 **73** 19.8	94 **78** 17.0
$J = 18$	86 **75** 12.8	86 **78** 9.3
$J = 20$	95 **79** 16.8	87 **78** 10.3

5 Conclusion

Multi-class classification is a fundamental task in machine learning. This paper presents the *abc-logitboost* algorithm and demonstrates its considerable improvements over *logitboost* and *abc-mart* on a variety of datasets.

There is one interesting UCI dataset named *Poker*, with 25K training samples and 1 million testing samples. Our experiments showed that *abc-boost* could achieve an accuracy $> 90\%$ (i.e., the error rate $< 10\%$). Interestingly, using LibSVM, an accuracy of about 60% was obtained[3]. We will report the results in a separate paper.

References

[1] Alan Agresti. *Categorical Data Analysis*. John Wiley & Sons, Inc., Hoboken, NJ, second edition, 2002.

[2] Peter Bartlett, Yoav Freund, Wee Sun Lee, and Robert E. Schapire. Boosting the margin: a new explanation for the effectiveness of voting methods. *The Annals of Statistics*, 26(5):1651–1686, 1998.

[3] Chih-Jen Lin. Private communications in May 2009 and August 2009.

[3] Colin B. Begg and Robert Gray. Calculation of polychotomous logistic regression parameters using individualized regressions. *Biometrika*, 71(1):11–18, 1984.

[4] Yoav Freund. Boosting a weak learning algorithm by majority. *Inf. Comput.*, 121(2):256–285, 1995.

[5] Yoav Freund and Robert E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *J. Comput. Syst. Sci.*, 55(1):119–139, 1997.

[6] Jerome H. Friedman. Greedy function approximation: A gradient boosting machine. *The Annals of Statistics*, 29(5):1189–1232, 2001.

[7] Jerome H. Friedman, Trevor J. Hastie, and Robert Tibshirani. Additive logistic regression: a statistical view of boosting. *The Annals of Statistics*, 28(2):337–407, 2000.

[8] Jerome H. Friedman, Trevor J. Hastie, and Robert Tibshirani. Response to evidence contrary to the statistical view of boosting. *Journal of Machine Learning Research*, 9:175–180, 2008.

[9] Yoonkyung Lee, Yi Lin, and Grace Wahba. Multicategory support vector machines: Theory and application to the classification of microarray data and satellite radiance data. *Journal of the American Statistical Association*, 99(465):67–81, 2004.

[10] Ping Li. Abc-boost: Adaptive base class boost for multi-class classification. In *ICML*, Montreal, Canada, 2009.

[11] Ping Li. Robust logitboost. Technical report, Department of Statistical Science, Cornell University, 2009.

[12] Ping Li, Christopher J.C. Burges, and Qiang Wu. Mcrank: Learning to rank using classification and gradient boosting. In *NIPS*, Vancouver, BC, Canada, 2008.

[13] Liew Mason, Jonathan Baxter, Peter Bartlett, and Marcus Frean. Boosting algorithms as gradient descent. In *NIPS*, 2000.

[14] Robert Schapire. The strength of weak learnability. *Machine Learning*, 5(2):197–227, 1990.

[15] Robert E. Schapire and Yoram Singer. Improved boosting algorithms using confidence-rated predictions. *Machine Learning*, 37(3):297–336, 1999.

[16] Ambuj Tewari and Peter L. Bartlett. On the consistency of multiclass classification methods. *Journal of Machine Learning Research*, 8:1007–1025, 2007.

[17] Tong Zhang. Statistical analysis of some multi-category large margin classification methods. *Journal of Machine Learning Research*, 5:1225–1251, 2004.

[18] Hui Zou, Ji Zhu, and Trevor Hastie. New multicategory boosting algorithms based on multicategory fisher-consistent losses. *The Annals of Applied Statistics*, 2(4):1290–1306, 2008.